C. Critique of resource management

The SUMEX community continues to be an extremely supportive environment in which to do research on uses of artificial intelligence in clinical medicine. The community has two equally vital resources -- the people with knowledge and interest in AI and the facility on which AI system development can proceed. They are equally excellent as resources, helping hands when faced with problems, and friendly support for continued productive research. The availability of INTERLISP; of a facility on which routine data processing functions (eg. manipulating magnetic tapes and making long listings) can take place; and of message-sending among remote users are all vital functions for our project. SUMEX provides them in an environment which is friendly and reliable.

D. Needs and plans for other computational resources

The computation facility at PMC is currently the source of all of the data being used by the PUFF/VM project, and it will continue in this capacity. We expect to link the two machines using a simple telephone dial-up link, but this represents the only system increment to the computational facility of the collaborative project. As the AI techniques developed under PUFF/VM enter routine clinical use at PMC, we have the requirement for system support on which these programs can execute. To date, we have been able to use the existing software development facilities on the PMC PDP-11 for this purpose, but we hope to be able to use more powerful mini-computer based software development facilities in the future.

E. Recommendations for future community and resource development

We perceive the evolution of our AI capability as moving from a highly speculative development state, for which the interactive development capabilities of SUMEX are vital, to a more stable but still changing validation-and-evaluation state. Ultimately we foresee rather stable specification of a program for routine clinical use. Thus, we see the need to transfer our AI techniques from the SUMEX PDP-10 to a local host. For this transfer, a principal long-range need is for software systems that will allow us to run AI systems on a mini-computer after they have been developed on the more powerful SUMEX facility. If the validation of PUFF/VM in the PMC clinical setting shows the programs to be effective in health care, then we hope and expect to be able to provide the capability on a routine basis.

4.1.7 RUTGERS COMPUTERS IN BIOMEDICINE

Rutgers Research Resource - Computers in Biomedicine

Principal Investigator: Saul Amarel
Rutgers University, New Brunswick, New Jersey

I. SUMMARY OF RESEARCH PROGRAM

A) Goals and Approach

The fundamental objective of the Rutgers Resource is to develop a computer based framework for significant research in the biomedical sciences and for the application of research results to the solution of important problems in health care. The focal concept is to introduce advanced methods of computer science - particularly in artificial intelligence - into specific areas of biomedical inquiry. The computer is used as an integral part of the inquiry process, both for the development and organization of knowledge in a domain and for its utilization in problem solving and in processes of experimentation and theory formation.

The Resource community includes 51 researchers - 30 members, 7 associates, and 14 collaborators. Members are mainly located at Rutgers. Collaborators are located in several distant sites and they interact, via the SUMEX-AIM and RUTGERS-10 facilities, with Resource members on a variety of projects, ranging from system design/improvement to clinical data gathering and system testing. At present, collaborators are located at the Mt. Sinai School of Medicine, N.Y.; Washington University School of Medicine, St. Louis, Missouri; Johns Hopkins Medical Center, Baltimore, Maryland; Illinois Eye and Ear Infirmary, Chicago, Illinois; the University of Miami, Florida, and the University of Pennsylvania, Philadelphia, Pennsylvania.

Resource activities include research projects (collaborative research and core research), training/dissemination projects, and computing in support of pilot user projects. The research projects are organized in three main AREAS OF STUDY. The areas of study and the senior investigators in each of these are:

(1) Medical Modeling and Decision Making (C. Kulikowski).

(2) Modeling Belief Systems and Commonsense Reasoning (C. F. Schmidt and N. S. Sridharan).

(3) Artificial Intelligence: Representations, Reasoning, and System Development (S. Amarel).

The training/dissemination activities of the Rutgers Resource (coordinated by R. Smith) include sponsorship of the Annual AIM Workshop - whose main objective is to strengthen interactions between AIM investigators, to disseminate research methodologies and results, and to stimulate collaborations and imaginative resource sharing within the framework of AIM. The third AIM Workshop was held at Rutgers on July 5 to 8, 1977. The fourth AIM Workshop is scheduled for June 25 to 28, 1978.

The RUTGERS-10 computer is being used not only for support of local research projects and AIM Workshop activities; but also for Pilot User Projects in the AIM community - within the general framework of the national AIM project. Computing activities in the Resource are coordinated by S. Levy.

B) Medical Relevance; Collaborations

A unique and novel aspect of our work is the creation of a network of clinical investigators to collaborate on the testing and continued development of the computer programs needed to accomplish the tasks mentioned in Section C.1.a. During 1977, the ophthalmological network (ONET) of glaucoma investigators continued to grow and has established itself, with several significant collaborative research projects currently under way. The consultation program for glaucoma using the causal association network (CASNET) model developed within the Rutgers Resource, was jointly presented by the ONET members at the 1976 meeting of the Association for Research in Vision and Ophthalmology. The results of the panel discussion where the CASNET/Glaucoma program was pitted against a group of experts are described in the recently published book Discussions On Glaucoma by Lichter and Anderson. The incorporation into the consultation program of alternative expert opinions on subjects currently under debate has been an important aspect of this work. The SUMEX-AIM shared computer resource has been essential to the activities of ONET.

The knowledge base and the strategies of our CASNET glaucoma consultation system are being strengthened and refined continuously in the ONET environment. The system is now at a point where it is considered by leading ophthalmologists as "highly competent-to-expert" in several subspecialties of glaucoma.

In the past year work has begun in several new research areas: hematology, endocrinology and enzyme kinetics. In addition to the design of consultation systems, underlying problems of knowledge representation and inference processes have been studied.

Research collaborations include the five medical centers of ONET, the Rutgers Medical School, and the University of Pennsylvania.

In the area of Belief Systems, collaboration is continuing with Prof. Andrea Sedlak and her research group at the University of North Carolina at Chapel Hill. Also, interactions continued with a group at Univ. of Massachusetts on the design of a knowledge-based system to learn rules for interpreting actions.

Our close contacts with the Stanford projects on Heuristic Programming (Drs. Buchanan, Feigenbaum, Lederberg) are continuing. The AI orientation and approach of these Stanford projects are very similar to ours. Graduate students at Rutgers continued to contribute to the development of the AI handbook - a project led by Dr. Feigenbaum at Stanford and intended to provide a network-accessible encyclopedic coverage of the AI field for the AIM community and AIM guests.

C) Progress Summary

1. Areas of Study and Projects:

a) Medical Modeling and Decision-Making

Research activities during the past year have concentrated on the development of new consultation systems, and associated investigations into representations of knowledge, strategies of inference, and planning. The evaluation and testing of the CASNET/Glaucoma system by the members of the ONET (Ophthalmological Network) has continued. Several new application areas for consultation have been started: hematology, endocrinology, and most recently, rheumatology. The methodology of modeling in enzyme kinetics has been studied, and preliminary data base and research support programs developed.

The collaborative activities of ONET, comprising investigators at the Mt. Sinai School of Medicine, Johns Hopkins University Washington University, the University of Illinois at Chicago, and the University of Miami, have continued through the testing of the CASNET/Glaucoma Model. The time-sequenced data base facilities have been used both by ONET and by hemophilia investigators at the College of Medicine and Dentistry of New Jersey (Rutgers Medical School). Investigators at this institution have also collaborated on the development of disease models for consultation in endocrinology. Collaborative investigation at the University of Pennsylvania has begun on problems of enzyme kinetics modeling.

In the area of general methods and systems for medical reasoning, there has been research into the representation of anatomical physiological knowledge in the form of descriptive models that interact with clinical-level models for decision-making. The AIMDS descriptive system has been used to model the visual pathways and to implement the generalized physiological reasoning that permits the interpretation of patterns of visual field loss. A related project involves the 3-dimensional graphical representation of the 'hill of vision' for clinical interpretation.

The investigation of reasoning schemes for diagnosis and treatment has been carried out (and new representational elements included) in two systems currently under development: IRIS and XPERT. The former is a semantic network-based system for propagating inferences written in INTERLISP. The latter is a modular consultation system that is being designed for ease in knowledge acquisition from experts and for transferability to small machines (the PDP-11/60 is being used as a prototype).

Clinical investigations in hemophilia, thyroid disease and hypertension have been aided by Resource support and development of the BRIGHT system.

b) Modeling of Belief Systems and Commonsense Reasoning

The last year has been a period of intensive system development in AIMDS and one of intensive conceptual development in BELIEVER.

AIMDS is intended to be a representational framework, which provides a combination of description languages, a logical language and several major procedures for creating and matching large description structures, and facilitates the construction of a variety of knowledge-based systems. Heavy emphasis is given to the representation and modeling of actions which may have side effects. A context mechanism is available that makes it possible to store and process plans, maintain belief models, and answer "what if" questions about actions. The intent is to have the user describe the knowledge being applied as needed, and consequently are expected to be simple to build, test, and also to understand.

The BELIEVER system is an example of a psychological theory that uses the tools of artificial intelligence, to construct a psychological theory of the process of understanding observed actions of others. Understanding observed action sequences involves the attribution of a plan to the actor and identifying the actor's goals. A plan is attributed to the actor even if the plan is not executed to completion. The attributed plan must be a plan that would realize the given goal, with respect to the observer's model of the actor's beliefs.

The plan must correspond to the observed actions. The goal identified must be attributable to the actor, i.e., there must be a motivational reason derivable for the goal for the actor. The evidence from subject experiments of recall and summarization suggest strongly that the observers develop a model of the actor's beliefs and intentions and that they focus on one hypothesis about the actor's plan. If the subsequent actions match the expectations this serves to confirm parts of the plan and allows the observer to refine the plans used for monitoring the action sequence. If subsequent actions disconfirm the expectations, this requires revisions to the hypothesized plan structure (perhaps including the identified goal), attempting to retain as much of the original interpretation as is possible.

The principal activity in BELIEVER has been the further clarification given to the notion of the "Hypothesize and Revise" paradigm and its articulation for the Plan Recognition problem in the form of about a dozen hypothesis revision rules. The arguments for the paradigm and a description of the revision rules are given in a paper that is to appear.

The BELIEVER and AIMDS system are programmed in a combination of RUCILISP and FUZZY. All of the effort involved in the design, construction, and testing of the psychological theory are carried out on the RUTGERS-10.

c) Artificial Intelligence; Representations, Reasoning, and Systems Development

Our work in this area continues to be oriented to collaboration with investigators in other Resource projects and to study of basic AI problems that are related to Resource applications. The collaborations involve adaptation and augmentation of existing AI methods and techniques to handle specific key problems identified in the application projects.

The close collaboration with investigators in the Belief Systems area has resulted in further development of the AIMDS system for handling problems of action interpretation of the type encountered in the domain of the BELIEVER theory.

The AIMDS system was operational in the summer of 1976. A vastly revised and speeded-up version was running in 1977 The system development of AIMDS in the past year has been directed at (a) the production of a user manual; (b) the production of several tutorial sessions; (c) general user engineering - adding a number of user controlled options, easy specification of these options, and making several of the functions provide helpful text; (d) adding more power in the various description languages that the system provides the user; (e) extending the system to handle multiple contexts in the data base; (f) adding a feature whereby defaults can be computed for missing relations; (g) construction of a more powerful Match function for comparing the description of one object with another. The AIMDS system continues to be an invaluable component of the the research in the construction of the BELIEVER theory. The combination of description and logical languages provided by the system and the various major processes play a significant role in the way the theory is made explicit.

In an attempt to test the generality of the representation available in AIMDS, we have programmed a learning algorithm which learns structural models from examination of examples and near misses, in the way of Pat Winston. The learning program has certain novel features, which eliminates the need to store and process past training instances. We have introduced a separate memory model in which a summary description of the past training instances are maintained. The memory model used for the blocks world is adequate to duplicate the learning phenomena demonstrated by Winston, while operating with a reduced memory load. The memory model is flexible and may be redesigned for different applications quite easily by changing its specifications written in the description language provided by AIMDS. An effort is being made to investigate this form of learning in the domain of C13 NMR spectra. It is anticipated that this form of learning which involves concept refinement by examining examples and near-misses will become another component of the Believer process.

One project currently underway in the general area of <u>natural language interfaces</u> is concerned with the development of a method for the parsing and interpretation of medical notes, and for guiding the researcher in the use and extension of the system's capabilities. By enabling the system to generate and track a closely constrained set of syntactic and semantic expectations during the parsing process, we hope to be able to permit it to deal intelligently with some of the problems arising from syntactic and semantic ambiguity and the widespread ellipsis of medical notes. New schemes for organization of this tracking strategy are being explored.

This year, in a continuing project in <u>hypothesis formation</u>, we have made further progress in the computer acquisition of domain knowledge from empirical data. The domain knowledge is in the form of weighted production rules, and a set of production rules can be represented as a stochastic graph. We obtained a general result about loop-free interpretation of a stochastic graph under max and min operations. We also wrote three programs (one in SITBOL, the other two in FORTRAN) as a preliminary implementation of a system for the semi-automatic construction of a succinct graph model from a large data base, and for the use of a model in knowledge-directed evidence gathering and decision making once it has been constructed. We began the testing of this system of three programs on a data base of case histories of glaucoma patients.

In the area of <u>theory formation in programming</u> our main effort has been directed to representations of programs in various stages of specification, and to methods for acquiring knowledge and for managing a knowledge base for the theory formation system.

During this period work continued on the development of a supportive local programming environment for our research. The FUZZY AI language is now relatively stable. In addition to its use in the Resource, as an implementation tool and as a means of modeling approximate algorithms, FUZZY was exported to a number of other sites in the United States and Western Europe. The RUTGERS/UCI LISP system, on which FUZZY is built, has been improved considerably over the past year. A number of new I/O facilities have been added, including a random access capability, support for sub-file directories, lower-to-upper case character conversion, and user-definable ersatz device names. String storage utilization was improved, and several new string-handling functions were incorporated into the system. The program commenting facility was improved and made storage efficient, and several new source-file handling features were added. A number of new functions and macros were added, including a sort/merge package and an ALGOL-like DO macro. The use of complex macros in interpreted code was optimized via the automatic saving of macro expansions. In addition, the compiler was extensively debugged and its generated code optimized. RUTGERS/UCI LISP is now in use at approximately 25 sites in the United States and Western Europe, and continues to spread. Some of the more active sites (Carnegie-Mellon University, University of California at Irvine, and University of Texas) have actively collaborated with Rutgers in system development activities.

2. AIM Dissemination/Training; The AIM Workshop

The third annual AIM Workshop took place July 5 to 8, 1977, at the Continuing Education Center of Rutgers University, and about 50 invitees attended.

The program included review of recent developments at centers where AI research is being done, with particular emphasis to work in medicine, biochemistry, medical information systems, plus related areas such as vision, speech understanding, and AI in education. In addition, an open meeting of the AIM Advisory Committee took place, and several panel discussions and working groups discussed important technical and management issues in depth. Scheduled panels had as topics AIM Research Management and Knowledge Engineering. As in the past, computer facilities were available for hands-on experimentation and demonstration of AI systems running at various sites such as SUMEX-AIM, SRI, and Rutgers, accessed via the TYMNET and ARPANET.

This year, the Workshop dealt more intensely with a smaller set of issues put to a smaller group of workers in the field. Technical reports were given by those individuals currently most active in AI; the panel discussions were planned to center around research policy and management issues. This format of the Workshop was chosen since it is the view of many senior researchers and representatives of funding groups that AI itself has made much progress, that reasonable formalisms and practical programming tools for knowledge representation and inference are now available, that much experience has been gained in developing systems in Research and Development environments, and that further progress in the direction of "exporting" systems to practical user environments requires careful thought about management and engineering issues.

A feature of the Workshop that was viewed as being particularly successful was the dynamic scheduling of the working groups. Under the direction of Dr. C. Kulikowski, technical program chairman of the Workshop, five working groups met and discussed issues related to knowledge representation and communication and dissemination in more detail; the results of these working groups were reported to the entire workshop on the last day. This format stimulated a good deal of discussion and interaction.

Most of the sessions of the Workshop were videotaped by the Rutgers Instructional Television campus facility. These videotapes are available for viewing, and form a permanent record of the Workshop itself. In addition, the proceedings of the workshop has been prepared, as a summary of the Workshop, and is being distributed to the invitees.

A panel on Applications of AI was organized by S. Amarel from the IJCAI-77 Conference which was held in Boston, Mass. on August 22 to 26, 1977. This panel was intended to further supplement the dissemination activities of AIM by bringing to a wide audience of AI researchers a status report on current development in AI systems and by focusing on some of the key AI problems that were identified so far in AIM research.

3. Computing facilities; Pilot Projects

During this period several new Pilot Projects for the AIM community were initiated on the RUTGERS-10 computer. These include collaboration with SUMEX-AIM on the design and testing of MAINSAIL on our TOPS-10 system, and work on BRIGHT, a NIH-sponsored clinical data base project.

D) Up-To-Date List of Publications

Amarel, S. and C. Kulikowski (1972) "Medical Decision Making and Computer Modeling," Proc. of 5th International Conference on Systems Science, Honolulu, January 1972.

Amarel, S. (1974) "Inference of Programs from Sample Computations," Proc. of NATO Advanced Study Institute on Computer Oriented Learning Processes, 1974, Bonas, France.

Amarel, S. (1974) "Computer-Based Modeling and Interpretation in Medicine and Psychology: The Rutgers Research Resource," Proc. on Conference on the Computer as a Research Tool in the Life Sciences, June 1974, Aspen, by FASEB; also appears as Computers in Biomedicine TR-29. June 1974, Rutgers University, also in Computers in Life Sciences. W. Siler and D. Lindberg (eds.), Faseb and Plenum, 1975.

Amarel, S. (1976) Abstract of Panel on "AI Applications in Science and Medicine" in 1976 National Computer Conference Program, N.Y., June 7-10, 1976.

Amarel, S. (1976) "Summary of National Computer Conference panel on AI Applications in Science and Medicine" in SIGART Newsletter and SIGBIO Newsletter, Fall 1976.

Amarel, S. (1977) "Summary of Issues and Open Problems in AI Applications," in Report of panel on Applications of Artificial Intelligence, Amarel (ed.), in Proc. IJCAI-77, MIT, Cambridge, Aug. 1977; S. Amarel organized this panel session for IJCAI-77.

Bruce, B. (1972) "A Model for Temporal Reference and its Application in a Question Answering Program," in "Artificial Intelligence," Vol. 3, Spring 1972.

Bruce, B. (1973) "A Logic for Unknown Outcomes," Notre Dame Journal of Formal Logic; also appears as Computers in Biomedicine, TM-35, Nov. 1973, Rutgers University.

Bruce, B. (1973) "Case Structure Systems," Proc. 3rd International Joint Conference on Artificial Intelligence (IFCAI), August 1973.

Bruce, B. (1975) "Belief Systems and Language Understanding," Current Trends in the Language Sciences, Sedelow, and Sedelow (eds.) Houston, in press.

Chokhani, S. and C.A. Kulikowski (1973) "Process Control Model for the Regulation of Intraocular Pressure and Glaucoma," Proc. IEEE Systems, Man and Cybernetics Conf., Boston, November 1973.

Chokhani, S. (1975) "On the Interpretation of Biomathematical Models Within a Class of Decision-Making Procedures," Ph.D. Thesis, Rutgers University; also Computers in Biomedicine TR-43, May 1973.

Fabens, W. (1972) "PEDAGLOT. A Teaching Learning System for Programming Language," Proc. ACM Sigplan Symposium on Pedagogic Languages, January 1972.

Fabens, W. (1975) "PEDAGLOT and Understanding Natural Language Processing." Proc. of the 13th Annual Meeting of the Assoc. of Computational Linguistics, October 30 - November 1, 1975.

Kulikowski, C.A. and S. Weiss (1972) "Strategies for Data Base Utilization in Sequential Pattern Recognition," Proc. IEEE Conf. on Decision and Control, Symp. on Adaptive Processes, December 1972.

Kulikowski, C.A. and S. Weiss (1973) "An Interactive Facility for the Inferential Modeling of Disease," Proc. 7th Annual Princeton Conf. on Information Sciences and Systems, March 1973.

Kulikowski, C.A. (1973) "Theory Formation in Medicine: A Network Structure for Inference," Proc. International Conference on Systems Science, January 1973.

Kulikowski, C.A., S. Weiss and A. Safir (1973) "Glaucoma Diagnosis and Therapy by Computer," Proc. Annual Meeting of the Assoc. for Research in Vision and Ophthalmology, May 1973.

Kulikowski, C.A. (1973) "Medical Decision-Making and the Modeling of Disease," Proc. First International Conf. on Pattern Recognition, October 1973.

Kulikowski, C.A. (1974) "Computer-Based Medical Consultation - A Representation of Treatment Strategies," Proc. Hawaii International Conf. on Systems Science, Jan. 1974.

Kulikowski, C.A. (1974) "A System for Computer-Based Medical Consultation," Proc. National Computer Conference, Chicago, May 1974.

Kulikowski, C.A. and A. Safir (1975) "Computer-Based Systems Vision Care," Proceedings IEEE Intercon, April 1975.

Kulikowski, C.A. and Trigoboff, M. (1975) "A Multiple Hypothesis Selection System for Medical Decision-Making," Proc. 8th Hawaii International Conf. on Systems.

Kulikowski, C.A. and N.S. Sridharan (1975) "Report on the First Annual AIM Workshop on Artificial Intelligence in Medicine." Sigart Newsletter No. 55, December 1975.

Kulikowski, C.A. (1976) "Computer-Based Consultation Systems as a Teaching Tool in Higher Education," 3rd Annual N.J. Conf. on the use of Computers in Higher Education, March 1976.

Kulikowski, C.A., S. Weiss, A. Safir, et al (1976) "Glaucomal Diagnosis and Therapy by Computer: A Collaborative Network Approach" Proc. of ARVO, April 1976.

Kulikowski, C.A., S. Weiss, M. Trigoboff and A. Safir (1976) "Clinical Consultation and the Representation of Disease Processes: Some AI Approaches," AISB Conferences, Edinburgh, July 1976.

Kulikowski, C.A., S. Weiss, A. Safir, (1976) "The CASNET Consultation System for Glaucoma: Development by Collaborative Computer Networking," Biosciences Communication Seminars, Aspen, Colorado.

Kulikowski, C.A., and A. Safir (1976) "Concepts of Computer-Based Modeling in Optics and Refraction," Proceedings Association for Computing Machinery Conference, Houston, Texas, pp. 110-114.

Kulikowski, C.A., S. Weiss and A. Safir (1977) "Computer Networking in the Development of a Glaucoma Consultation System" EDUCOM Fall Conference.

Kulikowski, C.A. (1977) "Representation of Knowledge for Reasoning in Medical Consultation," Panel Presentation, IJCAI, Boston.

Kulikowski, C. A (1977) "Problems in the Design of Knowledge Bases for Medical Consultation," The First Annual Symposium in Medical Care, Washington, D. C.

LeFaivre, R. and A. Walker (1975) "Rutgers Research Resource on Computers in Biomedicine, H," Sigart Newsletter No. 54, October 1975.

LeFaivre, R. (1976) "Procedural Representation in a Fuzzy Problem-Solving System," Proc. Natl. Computer Conf., New York, June 1976.

LeFaivre, R. (1977) "Fuzzy Representation and Approximate Reasoning," submitted to IJCAI-77, MIT.

Mathew, R., C. Kulikowski, and K. Kaplan (1977) "Multilevel Representation for Knowledge Acquisition in Medical Consultation Systems," Proc. MEDINFO Conference, Toronto.

Mauriello, D. (1974) "Simulation of Interaction Between Populations in Freshwater Phytoplankton," Ph.D. Thesis, Rutgers University 1974.

Safir, A., S. Podos, C. Kulikowski, S. Weiss, et al (1976) "Glaucoma Consultation by Computer," Scientific Exhibit E-22, Ophthalmology Transactions, American Academy of Ophthl. and Otolar., Vol. 81, No. 4.

Schmidt, C. (1972) "A comparison of source unidimensional, multidimensional and set theoretic models for the prediction of judgements of trail implication," Proc. Eastern Psych. Assoc. Meeting, Boston, April 1972.

Schmidt, C.F. and J. D'Addamio (1973) "A Model of the Common Sense Theory of Intension and Personal Causation," Proc. of the 3rd IJCAI, August 1973.

Schmidt, C.F. and A. Sedlak (1973) "An Understanding of Social Episodes," Proc. of Symposium on Social Cognition, American Psych. Assoc. Convention, Montreal, August 1973.

Schmidt, C.F. (1975) "Understanding Human Action," Proc. Theoretical Issues in Natural Language Processing: An Interdisciplinary Workshop in Computational Linguistics, Psychology, Artificial Intelligence, Cambridge, Mass., June 1975. Also appears as Computers in Biomedicine, TM-47, June 1975, Rutgers University.

Schmidt, C. (1975) "Understanding Human Action: Recognizing the Motives," Cognition and Social Behavior, J.S. Carroll and J. Payne (eds.), New York: Lawrence Earlbaum Associates, in press. Also appears as Computers in Biomedicine, TR-45, June 1975, Rutgers University.

Schmidt, C.F., N.S. Sridharan and J.L. Goodson (1976) Recognizing plans and summarizing actions. Proceedings of the Artificial Intelligence and Simulation of Behavior Conference, University of Edinburgh, Scotland, July 1976.

Schmidt, C. (1976) Understanding human action: Recognizing the plans and motives of other persons. In (eds. J. Carroll and J. Payne) Cognition and Social Behavior, Potomac, Maryland: Lawrence Earlbaum Associates, 1976.

Schmidt, C.F. and J.L. Goodson (1976) "The Subjective Organization of Summaries of Action Sequences," 17th Annual Meeting of the Psychonomic Society, St. Louis, 1976.

Schmidt, C.F., N.S. Sridharan, and J. L. Goodson (1977) "Plan Recognition: A Hypothesize and Revise Paradigm" Proceedings of the 5th International Joint Conference on Artificial Intelligence, Cambridge, Mass.

Schmidt, C. F, N. S. Sridharan, and J. L. Goodson (1978) "The Plan Recognition Problem: An Intersection of Psychology and Artificial Intelligence," Artificial Intelligence, Special Issue on applications to the sciences and medicine, to appear.

Sridharan, N.S. (1976) "The Frame and Focus Problems in AI: Decision in Relation to the BELIEVER System. Proceedings of the Conference on Artificial Intelligence and the Simulation of Human Behavior, Edinburgh, July 1976.

Sridharan, N.S. (1976) "An Artificial Intelligence System to Model and Guide Organic Chemical Synthesis, Planning in Chemical Synthesis by Computer, American Chemical Society Press, September 1976.

Sridharan, N.S. and C.F. Schmidt (1977) "Knowledge-Directed Inference in BELIEVER" in Pattern-Directed Inference, D. Waterman and F. Hayes-Roth (Editors) Academic Press (Scheduled for September).

Sridharan, N.S. and F. Hawrusik (1977) "Representation of Actions with Side-effects". In Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Boston.

Srinivasan, C.V. (1973) "The Architecture of a Coherent Information System: A General Problem Solving System," Proc. of the 3rd IJCAI, August 1973.

Trigoboff, M. (1976) Propagation of Information in a Semantic Net," Proc. of the Conference on Artificial Intelligence and the Simulation of Behaviour, Edinburgh, Scotland, July 1976; updated version appears in CBM-TM-57, Dept. of Computer Science, Rutgers University, 1977.

Trigoboff, M. and C. Kulikowski (1977) "IRIS: A System for Propagation of Inference in a Semantic Network," Proc. Fifth International Joint Conference on Artificial Intelligence, Boston.

Tucker, S.S. (1974) Cobalt Kinetics in Aquatic Microcosms," Ph.D. thesis, Rutgers University.

Van der Mude, A. and A. Walker (1976) "Some Results on the Inference of Stochastic Grammars," abstract in Proc. Symposium on New Directions and Recent Results in Algorithms and Complexity. Dept. of Computer Science, Carnegie-Mellon University; Information and Control, (in press).

Van der Mude, A. and A. Walker (1978) "On the Inference of Stochastic Regular Grammars," Information and Control, (to appear).

Vichnevetsky, R. (1973) "Physical Criteria in the Evaluation of Computer Methods for Partial Differential Equations," Proc. 7th International AICA Congress, Prague, Sept. 1973; reprinted in Proc. of AICA, Vol. XVI, No. 1, Jan. 1974, European Academic Press, Brussels, Belgium.

Vichnevetsky, R., K.W. Tu, and J.A. Steen (1974), "Quantitative Error Analysis of Numerical Methods for Partial Differential Equations," Proc. 8th Annual Princeton Conference on Information Science and Systems, Princeton University, March 1974.

Walker, A. (1975) "Formal Grammars and the Regeneration Capability of Biological Systems," Journal Comp. and Syst. Sciences, Vol. 11, No. 2, 252-261.

Weiss, S. (1974) "A System for Model-Based Computer-Aided Diagnosis and Therapy," Parts I and II, Ph.D. Thesis, Rutgers University; also Computers in Biomedicine TR-27, Feb. 1974.

Weiss, S. (1976) "A System for Interactive Analysis of a Time-Sequenced Ophthalmological Data Base, Proceedings of Third Illinois Conference on Medical Information Systems.

Weiss, S., C. Kulikowski, and A. Safir (1978) "Glaucoma Consultation by Computer," Computers in Biology and Medicine, (in press).

Weiss, S., C. Kulikowski, A. Safir, "A Model-base for Computer-Aided Medical Decision-Making," Journal of Artificial Intelligence, (in press).

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

During the past year we have continued to use the SUMEX-AIM resource for program development and testing, for communications among collaborators distributed in different parts of the country and for preparation and running of the AIM Workshop.

Among new projects developed in recent months are the TENEX SAIL system, which will also be used under the new 2050T system at Rutgers. The SUMEX-AIM system was also used to rewrite BRIGHT to work on TENEX, to debug TOPS-20 software, and to further develop the IRIS system. In another application, SUMEX-AIM was instrumental as a communications vehicle when the conversion of the Rutgers facility to the new 2050T system was being planned.

We continue to access SUMEX-AIM via TYMNET, and to a smaller extent via ARPANET. SUMEX-AIM played a key role in consolidating our network of collaborators in ophthalmology (ONET) and in providing the support needed for establishing a productive collaboration among the ONET investigators. It has also been most useful in communicating, planning and helping to set up the information pool for the third AIM Workshop.

Computing in the Rutgers Research Resource continues to be distributed between SUMEX-AIM and the RUTGERS-10. The two computers are providing complementary resources for our research and for our national collaborations. At present, the distribution of our computing is about 20 to 1 between RUTGERS-10 and SUMEX-AIM. Most of our work on SUMEX-AIM is done in INTERLISP (about 80% of our total connect hours) and the rest devoted mainly to communications and to limited program testing within ONET.

The SUMEX-AIM facility was used for demonstrations of AIM programs in first year classes and in second year seminars at the Rutgers Medical School: CASNET, MYCIN, INTERNIST and PARRY were accessed interactively in these classes and seminars. Another innovative use of SUMEX-AIM has been the collaborative development of the AI HANDBOOK, which is intended to provide a computer-based, network-accessible encyclopedia on Artificial Intelligence for the AIM community and guests. The AI HANDBOOK, was initiated by Dr. E. Feigenbaum and his students at Stanford. For the second year, a Rutgers graduate class, given by Dr. S. Amarel, worked on the Handbook and contributed several articles.

Finally, the SUMEX-AIM bulletin board continues to play an important role in communicating ideas and information on services among users. Since the MYCIN group at Stanford regularly posts summaries of meetings and other technical information on the MYCIN bulletin board, we have been able to keep track of their program and problems. This was particularly useful for our work on IRIS, where concepts close to the MYCIN CF formalism are being studied.

III. RESEARCH PLANS (8/78 - 7/81)

A. Long Range Project Goals and Plans

We are planning to continue along the main lines of research that we have established in the Resource to date, with emphasis on broadening our activities in the medical systems area. We are also planning to increase our participation in AIM dissemination and training activities, and to enhance the RUTGERS-10 computer in order to provide the support needed within the Resource and to increase the shared computing capabilities of the national AIM community.

B. Plans for Computing at Rutgers and at Stanford; development of the RUTGERS-10 facility as a node in the AIM network

By combining the NIH-recommended funding for computing over the renewal period of the Rutgers Resource (Dec. 1, 1977 to Nov. 30, 1980) and the Rutgers commitment for support of our research computing, the following plan for computing was developed and is now being implemented.

1. The current KI system at Rutgers was completely purchased in Dec. 1977. A new configuration was ordered to replace the present system in July 1978. The new system will be a KL-2050T with 512 K words of core and a TOPS-20 operating system.

2. The RUTGERS-10 computer is now controlled by the newly created Laboratory for Computer Science Research (LCSR) in which the Research Resource is administratively located; it is operated by a special CCIS group under a facilities management arrangement with LCSR. As Director of LCSR and PI of the Resource, Dr. S. Amarel has policy responsibility for the RUTGERS-10 facility vis-a-vis BRP and the University.

3. The user capacity of the RUTGERS-10 will be allocated as follows: (a) 55% to NIH grant activities - of this share, 3/4 will be allocated within the Resource for collaborative and core research, and 1/4 for the national

AIM community. [Thus, about 14% of user capacity of the RUTGERS-10 will be allocated to the AIM community.]; (b) 45% to Rutgers - this will be mainly devoted to computer science research (outside the NIH-supported Resource) and to advanced instruction in DCS. The user capacity is assumed to be 80% of total capacity, with the remaining 20% devoted to local systems activities and operations. Thus, an annual allocation of about 6,000 connect hours of KL (with about 1:60 compute to connect ratio) will be available to the AIM community starting in July 1978. About 20% of this will be needed for the AIM Workshop; the remaining suballocations will be devoted to national AIM users outside Rutgers. Based on projections of computing demand within the Rutgers Resource, the 1978 level of allocation may not be maintained in 2-3 years without enhancement of the KL configuration, although it is expected that at least 3,000 connect hours will be available in 1980.

4. The share devoted to the AIM community will be governed within the management framework of the existing national AIM committee structure - in the same manner that the national share of the SUMEX-AIM facility at Stanford is currently being governed. We propose to place special emphasis in the Rutgers AIM node on collaborative developments of knowledge-based systems in medicine. Also, since the TOPS-20 system (which represents the new line of DEC-supported operating systems) will be available on the Rutgers KL-2050T, it can provide the basis for technical experiments by SUMEX-AIM systems staff in preparation for future system changes in the AIM network. Such system support/planning activities for the AIM community should receive high priority in AIM usage on the RUTGERS-10.

This planning is based on the assumption that computing by investigators in the Rutgers Resource will continue to be distributed between the RUTGERS-10 and SUMEX-AIM, with the bulk of computing being done at Rutgers. We expect a demand level of about 2,000 connect hours per year on SUMEX-AIM - mainly for INTERLISP-based system developments, communications, special software developments and collaborative activities including AIM Workshops and the AI handbook.

E. FUNDING INFORMATION

1. Title of Grant or Contract:

   Rutgers Research Resource on Computers in Biomedicine

2. Principal Investigator(s): Saul Amarel
   Title and Institutional Affiliation: Professor and Chairman
   Department of Computer Science; and
   Director, Laboratory for Computer Science Research
   Rutgers University

3. Funding Agency: National Institutes of Health
   Biotechnology Resources Program
   Division of Research Resources

4. Grant or Contract Identification Number(s):

No. 2 P41 RR-00643

5. Total award term(s):

Three year renewal grant

Dates	Funding Amounts (DIRECT COSTS ONLY)
12/1/78-11/30/78	$505,823
12/1/78-11/30/79	522,005
12/1/78-11/30/80	437,064
	$1,464,892

6. Current Term(s), Dates, and Funding Amount(s):

12/1/77-11/30/78 $505,823

4.1.8 SIMULATION OF COMPREHENSION PROCESSES

Simulation of Comprehension Processes (SCP)

James G. Greeno and Alan M. Lesgold
Learning Research And Development Center
University of Pittsburgh

I. SUMMARY OF RESEARCH PROGRAM

The SCP project has only been on SUMEX for several weeks, so there is not too much to report yet. Even worse, these weeks have been mostly the period in which various professional meetings and other responsibilities have kept the investigators (Greeno and Lesgold) away from their work. Nonetheless, there is some progress to report and other information to be provided for this document.

Technical Goals

The goals of this project remain the simulation of young children's behavior in arithmetic and reading tasks, done in such a way that various levels of proficiency, from moderate dysfunction to expertise, can be modeled. In the reading work, the emphasis is on the interaction of several components of the reading process: word recognition, sentence parsing, and relating new sentence content to what has already been understood. The project begins from the empirical basis of a large number of studies showing that various word processing skills are less developed in deficient readers. Lesgold and Perfetti (in press) have discussed this evidence and have suggested that poorly developed word processing skills use too much of a limited processing capacity, thereby indirectly hampering comprehension skills that may otherwise be operating at acceptable levels.

Regrettably, the evidence to support this point of view has been largely correlational. To support a more detailed and empirically verified specification of the sources of reading dysfunction, we have adopted the strategy of combining several converging empirical investigations with a simulation of the reading process that can more clearly specify the convergence of evidence. Thus, we are conducting a number of empirical studies of comprehension processes, have reviewed a large body of evidence on word recognition (e.g., Perfetti & Lesgold, in press; Lesgold & Perfetti, in press, 1977), and are conducting a longitudinal study of changes in children's word processing and reading abilities throughout the course of primary-grades reading instruction.

In our work on arithmetic, we are continuing the development of a model of comprehension processes related to elementary arithmetic concepts. The model simulates a process of solving simple word problems, constructing an integrated representation of a problem using schematic knowledge about quantitative relationships (Greeno, 1978). We are conducting a developmental study to determine the degree to which semantic and linguistic factors, rather than arithmetic knowledge, are responsible for children's difficulty in solving problems at early grade levels.

Medical Relevance and Collaboration

The range of ability levels we will be dealing with in arithmetic and reading includes children who are below average and, in the case of reading, some children who are classed as "learning-disabled." By providing a framework within which the effects of differing levels of skill acquisition can be understood, we hope to eliminate the spurious use of vague medical categories such as "minimal brain damage," etc., and thus more clearly delimit the cases in which there is a real medical problem in a child whose achievement in math and reading is poor. Toward this end, we have collaborated with Isabel Beck of the Learning Research and Development Center in doing comparisons of our data from "normal" children in reading with her data on "LD" children. The comparisons should, at a later stage of the simulation work, play a role in the design of an overall model for reading.

Progress Summary

The few weeks that we have been on SUMEX have been spent in learning the system and in interacting with the ACT project (see below) to develop a complete understanding of the resources that can be borrowed from them rather than reinvented. The empirical progress on the reading work is described in the publications listed below and in a paper presented on March 30, 1978, at the annual meeting of the American Educational Research Association. Recent progress on the arithmetic work will be reported in papers to be presented at the May meetings of the Midwestern Psychological Association, as well as in Greeno's paper cited below.

List of Relevant Publications

Greeno, J.G. Some examples of cognitive task analysis with instructional implications. Paper presented at the ONR/NPRDC Conference on Aptitudes, Cognition, and Instruction, San Diego, March, 1978.

Greeno, J.G. A study of problem solving. In R. Glaser (Ed.), Advances in instructional psychology. Hillsdale, NJ: Erlbaum, in press.

Heller, J.I., & Greeno, J.G. Semantic processing in arithmetic word problem solving. Paper presented at Midwestern Psychological Association, May, 1978.

Lesgold, A.M., and Perfetti, C.A. Interactive processes in reading. Discourse Processes, in press.

Perfetti, C.A., & Lesgold, A.M. Discourse processing and sources of individual differences. In P. Carpenter & M. A. Just (Eds.), Cognitive processes in comprehension, Hillsdale, NJ: Erlbaum, 1977.

Perfetti, C.A., & Lesgold, A.M. Coding and comprehension in skilled reading. In L.B. Resnick & P. Weaver (Eds.), Theory and practice of early reading, Hillsdale, NJ: Erlbaum, in press.

Riley, M.S., & Greeno, J.G. Importance of semantic structure in difficulty of arithmetic word problems. Paper presented at Midwestern Psychological Association, May, 1978.

Funding support status

A. ONR
1. Analysis of Formal and Informal Reasoning in Problem Solving
2. James Greeno, Research Associate and Professor of Psychology, Learning Research and Development Center, University of Pittsburgh..
3. Office of Naval Research
4. Contract Number N0014-78-C-0022
5. October 1, 1977 through September 30, 1980; $198,000.
6. October 1, 1977 through September 30, 1978; $62,616.

B. NIE
1. Adaptive Education
2. Robert Glaser (James Greeno and Alan Lesgold are both Research Associates of the Learning Research and Development Center and are funded out of this grant).
3. National Institute of Education
4. Grant Number OB-NIE-78-0115.
5. December 1, 1977 through November 30, 1978; $3,800,000.
6. December 1, 1977 through May 30, 1978; $700,000.

(N.B. the current level of funds expended on the specific research projects of Greeno and Lesgold from this funding source are $70,140 and $55,760 per annum, direct costs, respectively. In addition, Greeno and Lesgold use substantial computer resources and other technical resources for experimentation and data analysis that are funded from the overall grant but not charged to their specific budgets.)

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

Collaborations and medical use of programs via SUMEX

No such collaborations have yet been implemented.

Sharing and interactions with other SUMEX-AIM projects

We have been in continual touch with the members of the ACT project, primarily because our simulations will be developed within the ACT system. However, these interactions have been bilateral in effect. As we have asked questions, details of ACT have become better understood and better implemented.

Critique of resource management

We have enjoyed clearly adequate resources so far. One suggestion to be made is that national users be regularly informed of the telecommunications costs associated with access to SUMEX-AIM so that responsible decisions can be made regarding use of the system.

III. RESEARCH PLANS

Long range project goal and plans

As work on modeling arithmetic knowledge proceeds, we hope to be able to develop theoretical analyses of the process of acquisition of elementary mathematics. In this work, we anticipate further collaboration with the ACT project, in which recent efforts have been focused on mechanisms of learning. Preliminary discussions of collaborative work on the acquisition of strategic knowledge in geometry problem solving, such as that described in Greeno (in press).

The goals for reading are similar, but tend to more heavily emphasize development of models of components of the reading process that are able to easily characterize the wide range of individual differences in reading that are found in the U.S. The work will involve continual alternation between empirical investigations and use of SUMEX-AIM for modeling. We also anticipate substantial interactions with Anderson and Kline of the ACT project to continue over the life of the project.

Justification and requirements for continued SUMEX use

The primary justification for our access to SUMEX-AIM is that it is the "home" of both the versions of ACT that we need for our work and the ACT researchers with whom we can benefit from continual interaction, exchange of developing simulation programs, etc. An additional important justification is that we do not presently have access to a computer system that can support ACT's more useful versions, nor is such a facility likely to become available. A final benefit is our ability to interact with other researchers through ARPAnet. We anticipate that other workers will be able to run demonstrations of our simulations, as they develop, through the SUMEX-AIM guest arrangements.

At the present time, our allocations of disk space within SUMEX are adequate, though it is conceivable that before a year passes we will need a small increment in space. If so, a request will be made to the Executive Committee at that time.

Needs and plans for other computer resources beyond SUMEX-AIM

We make substantial use of the computer systems at the University of Pittsburgh for all the data analysis and on-line experimental control services that the empirical side of our research requires. In addition, a more primitive version of ACT has been installed on the local system here. Thus, our use of

SUMEX-AIM is limited to the specific simulation needs which cannot otherwise be accomplished.

Recommendations for future community and resource development

We do not have enough experience on SUMEX-AIM to make recommendations yet.

4.2 STANFORD PROJECTS

The following group of projects is formally approved for access to the Stanford aliquot of the SUMEX-AIM resource. Their access is based on review by the Stanford Advisory Group and approval by Professor Lederberg as Principal Investigator. As noted previously, the DENDRAL project was the historical core application of SUMEX. Although this is described as a "Stanford project," a significant part of the development effort and of the computer usage is dedicated to national collaborator-users of the DENDRAL programs.

4.2.1 AIHANDBOOK PROJECT

Artificial Intelligence Handbook

E. A. Feigenbaum
Stanford Computer Science Department

I. SUMMARY OF RESEARCH PROGRAM

A. Technical Goals

The AI Handbook is a compendium of knowledge about the field of Artificial Intelligence. It is being compiled by students and investigators at several research facilities across the nation. The scope of the work is broad: Two hundred articles cover all of the important ideas, techniques, and systems developed during 20 years of research in AI. Each article, roughly four pages long, is a description written to be suitable both for the non-AI specialist and students of AI. Additional articles serve as Overviews, which discuss the various attempts within a subfield, the issues, and the problems.

There is no comparable resource for AI researchers and other scientists who need access to descriptions of AI techniques like problem solving or parsing. The research literature in AI is not generally accessible to outsiders. And the few textbooks that exist are not nearly broad enough in scope to be useful to a scientist working primarily in another discipline who wants to do something requiring knowledge of AI. Furthermore, we feel that some of the Overview articles are the best critical discussions available anywhere of activity in the field.

To indicate the scope of the Handbook, we have included an outline of the articles as an appendix to this report (see Appendix I on page 217).

B. Medical Relevance and Collaboration

The AI Handbook Project was undertaken as a core activity by SUMEX in the spirit of community building that is the fundamental concern of the facility. We feel that the organization and propagation of this kind of information to the AIM community, as well as to other fields where AI is being applied, is a valuable service that we are uniquely qualified to support.

C. Progress Summary

Because our objective is to develop a comprehensive and up-to-date survey of the field, our article-writing procedure is suitably involved. First drafts of Articles are reviewed by the staff and returned to the author (either an AI scientist or a student in the area). His final draft is then incorporated into a Chapter, which when completed is sent out for review to one or two experts in that particular area, to check for mistakes and omissions. After corrections and comments from our reviewers are incorporated by the staff, the manuscript is edited, and a final computer-prepared, photo-ready copy of the Chapter is generated.

We expect the Handbook to reach a size of approximately 700 pages. By that estimate, we are now 1/4 finished and expect to reach the halfway point during the Summer of 1978, since a good deal of the material outstanding has already been submitted in first draft form. Volume I of the Handbook, which will include the material completed in Summer 1978, will cover AI research in Heuristic Search, Representation of Knowledge, AI Languages, Natural Language, Speech Understanding, Applications-oriented AI research, and Automatic Programming. It may also include the chapters on Information Processing Psychology, Learning and Planning. Researchers at Stanford University, Rutgers University, SRI International, XEROX PARC, RAND Corporation, and others have already contributed articles or reviewed chapters.

D. List of Relevant Publications

Volume I of the AI Handbook will appear in preliminary from as a Stanford Computer Science Technical Report in the Fall of 1978.

E. Funding Support Status

The Handbook Project is currently supported under the Heuristic Programming Project contract with the Advance Research Projects Agency of the DOD, contract number MDA 903-77-C-0322, E. A. Feigenbaum, Principle Investigator. During the next year, SUMEX core research funds will provide partial personnel support for approximately 1.4 FTE.

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

A. Collaborations and medical use of programs via SUMEX

We have had a modest level of collaboration with a group of students and staff at the Rutgers resource, as well as occasional collaboration with individuals at other ARPA net sites.

B. Sharing and interactions with other SUMEX-AIM projects.

As described above, we have had moderate levels of interaction with other members of the SUMEX-AIM community, in the form of writing and reviewing Handbook material. During the development of this material, limited arrangements have been made for sharing the emerging text. As final manuscripts are produced, they will be made available to the SUMEX-AIM community both as on-line files and in the hardcopy, published edition.

C. Critique of Resource Management

Our requests of the SUMEX management and systems staff, requests for additional file space, directories, systems support, or program changes, have been answered promptly, courteously and competently, on every occasion.

III. RESEARCH PLANS (8/78 - 7/81)

A. Long Range Project Goals

The following is our tentative schedule for completion and publication of the AI Handbook:

May through August, 1978 - Various Chapters of Volume I will be available as Stanford Computer Science Technical Reports, distributed to the AI community for peer review.

Fall 1978 - After revision reflecting criticism of the reviewers, material for Volume I will be submitted for publication.

Fall 1978 through Spring 1979 - Development and refinement of material in remaining Chapters, Volume II.

Summer 1979 - Completion and publication of Volume II.

B. Justifications and requirements for continued SUMEX use

The AI Handbook Project is a good example of community collaboration using the SUMEX-AIM communication facilities to prepare, review, and disseminate this reference work on AI techniques. The Handbook articles currently exist as computer files at the SUMEX facility. All of our authors and reviewers have access to these files via the network facilities and use the document-editing and formatting programs available at SUMEX. This relatively small investment of resources will result in what we feel will be a seminal publication in the field of AI, of particular value to researchers, like those in the AIM community, who want quick access to AI ideas and techniques for application in other areas.

C. Your needs and plans for other computational resources

We use document preparation facilities (the XEROX Graphics Printer) at the Stanford AI Laboratory.

D. Recommendations for future community and resource development

None.